Session 5 Overview

Processors

Chair: Héctor Sánchez, Freescale Semiconductor, Austin, TX

Associate Chair: Georgios Konstadinidis, Sun Microsystems, Sunnyvale, CA

Technology, as we have known it during the golden years of Moore's law, is facing significant challenges related to transistor and metallization scaling, as well as power and thermal control. High-performance transistors exhibit very high leakage and variability. Interconnect delay and power are hurdles for continued single-core processing performance improvements. Dynamic power dissipation and heat removal limit the industry's ability to continue the unprecedented performance gains of the last 20 years. Supply voltage scaling has come to a halt. Technology options used to optimize the power-performance trade-off are running out of steam. Variable voltage and frequency single-core processing can only partially alleviate the situation.

Given all these issues, is it time for a *fundamental* shift in processor computing implementation philosophy? Maybe so. A look at the processor papers this year reveals that all of the major companies involved in high-performance computing are rethinking the way to deliver improved power/performance products. No high-performance single-core processor papers are disclosed, a first for ISSCC. However, there are five multi-core processor innovations by five different companies. How is this new paradigm-shift to deliver Moore's law on performance/power scaling without transistor scaling? The answer lies in the combination of architectural changes, technology improvements, and extensive use of power/performance-efficient multi-core designs. On the architectural side, the increased parallelism allowed by multi-threading alleviates the requirement for high-frequency processing with a consequence of more power-efficient designs. This in turn, reduces the technology requirements for transistor scaling, resulting in less power-hungry transistors. Aggregate computing gains are accomplished by replication of simpler multi-threaded cores whose overall power/performance/cost is substantially superior to single core power/thermal-limited designs.

In Paper 5.1, the first generation of the power-efficient 64b-Niagara SPARC processor comprising 8 cores able to simultaneously handle 32 threads is described. It includes a high-bandwidth crossbar, a 3MB L2 cache and four DDR2 interfaces. The 378mm^2 die implemented in a 90nm technology consumes 63W at 1.2GHz. In Paper 5.2, a multi-core RISC processor with a number of security engines and network function accelerators results in a high-performance power-efficient system on a chip. It dissipates 25W at 600MHz and is fabricated in a 1.2V 0.13μm CMOS process with 9 layers of copper interconnects.

The first reported processor implementation in a 65nm technology is presented in Paper 5.3. This 435mm^2 dual-core 64b Xeon™ MP processor has 1.328B transistors, operates in excess of 3GHz and executes two threads per core. The chip contains a unified 1MB L2 cache per core and a shared 16-way set-associative unified L3 cache. A microprocessor featuring 2 Hammer cores, an on-chip DDR2 memory controller, and the Pacifica architectural support for virtualization is presented in Paper 5.4. It is processed in a 90nm triple-V_t partially-depleted SOI technology with 9 layers of copper interconnects and operates at 2.6GHz at 1.35V with power dissipation of 95W. Paper 5.5 showcases a 146mm^2 dual-core Power™ architecture 64b processor created using a 90nm dual strained-silicon SOI technology with operation up to 3GHz. The dual-core processor has split clock domains and power planes, a 1MB L2 cache per core, and a shared processor interconnect bus.

The fastest integer execution unit reported to date, part of a 4th-generation Intel Pentium® 4 processor, is presented in Paper 5.7. The 9GHz operation at 1.3V 70°C is enabled by the use of 65nm technology and circuit design optimization for low latency and power.

The need for high aggregate bandwidth is satisfied by a 170GB/s crossbar for a multiprocessor server realized with only 10 LSIs as described in Paper 5.6. A 1.333Gb/s single-ended signal transmission with a driver using pre-emphasis and a receiver using a novel data-synchronous scheme is discussed. The total bandwidth of the address crossbar LSI is 1.23Tb/s with 704 drivers and 352 receivers.

 1-4244-0079-1/06/$20.00

5.1 A Power-Efficient High-Throughput 32-Thread SPARC Processor **1:30 PM**
A. Leon, Sun Microsystems, Sunnyvale, CA

The first generation of Niagara SPARC processors implements a power-efficient multi-threading architecture to achieve high throughput with minimum hardware complexity. The design combines eight 4-threaded 64b cores, a high-bandwidth crossbar, a shared 3MB L2 Cache and four DDR2 DRAM interfaces. The 90nm 378mm^2 die consumes 63W at 1.2GHz. Memory design techniques to support the high bandwidth are also discussed.

5.2 A 16-Core RISC Microprocessor with Network Extensions **2:00 PM**
V. Yalala, Cavium Networks, Marlboro, MA

A multi-core RISC processor is integrated with a number of security engines and network function accelerators creating a high-performance power-efficient SoC. It contains 180M transistors, dissipates 25W at 600MHz and is fabricated in a 1.2V 0.13µm CMOS process with 9 layers of copper interconnect using FSG dielectric and C4 bumps.

5.3 A Dual-Core Multi-Threaded Xeon™ Processor with 16MB L3 Cache **2:30 PM**
S. Rusu, Intel, Santa Clara, CA

A dual-core 64b Xeon™ MP processor is implemented in a 65nm 8M process. The 435mm^2 die has 1.328B transistors. Each core has two threads and a unified 1MB L2 cache. The 16MB unified, 16-way set-associative L3 cache implements both sleep and shut-off leakage reduction modes.

5.4 A 2.6GHz Dual-Core 64b x86 Microprocessor with DDR2 Memory Support **3:15 PM**
M. Golden, Advanced Micro Devices, Sunnyvale, CA

A microprocessor featuring 2 Hammer cores and an on-chip DDR2 memory controller implements Pacifica architectural support for virtualization. It is fabricated in a 90nm triple-V_t partially-depleted SOI process with 9 layers of copper interconnect. The chip achieves a clock frequency of 2.6GHz at 1.35V while dissipating 95W.

5.5 A 64b CPU Pair: Dual- and Single-Processor Chips **3:45 PM**
E. Cohen, IBM, Essex Junction, VT

Two Power™-architecture 64b microprocessor chips are fabricated in 90nm dual strained-silicon SOI technology. The dual-processor chip has split clock domains and power planes, 1MB L2 cache per core and a shared processor interconnect bus. The single-processor chip shares the dual's basic core and cache design.

5.6 High-Speed Interconnect for a Multiprocessor Server Using Over 1Tb/s Crossbar **4:15 PM**
J. Yamada, Fujitsu, Kawasaki, Japan

A 170GB/s crossbar for a multiprocessor server is realized with 10 LSIs. High density and low power are achieved with a 1.333Gb/s single-ended signal transmission, a driver using pre-emphasis, and a receiver using a data-synchronous scheme. The total bandwidth of the address crossbar LSI is 1.23Tb/s with 704 drivers and 352 receivers.

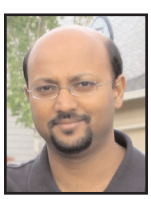

5.7 A 9GHz 65nm Intel Pentium®4 Processor Integer Execution Core **4:45 PM**
S. Wijeratne, Intel, Hillsboro, OR

In a 4th-generation 65nm Intel Pentium®4 processor, the previously low voltage swing, 2x microprocessor frequency, AGU and ALUs are replaced with domino-logic-based architectures optimized for low latency, 2x frequency, and lower power. These redesigned AGU/ALUs reduce normalized dynamic power by 50% over the previous generation, and together with the similarly optimized integer register file, enable a 9GHz 64b integer execution core at 1.3V and 70°C.